Experiential Knowledge and Engineering Education

William J. Clancey

Department of Computer Science
Knowledge Systems Laboratory
701 Welch Road, Building C
Palo Alto, CA 94304

and

Institute for Research on Learning
3333 Coyote Hill Road
Palo Alto, CA 94304

Presented at the McDonnell "AI and Engineering Education" Workshop, Asilomar, CA, December 13–15, 1987.

{Adapted from commentary entitled, "Functional principles and situated problem solving" for John Anderson's article "Methodologies for studying human knowledge," to appear in Behavioral and Brain Sciences, 1987.}

Current expert system research involves design of "generic tools." These programs formalize recurrent, domain-general knowledge organization and inference procedures for particular tasks (e.g., diagnosis, planning, control) in multiple domains (e.g., medicine, electronics). The formalization and study of this level of knowledge can be contrasted with traditional psychology by considering how AI researchers use the word "task"--a *kind of problem* (such as diagnosis or programming), not a specific problem to solve (patient to diagnose or program to write).

Much research on application of AI to education focuses on relatively formal problems, such as geometry and Lisp programming. From the perspective of expert systems developed for scientific and engineering problems, cognitive science research in mathematics, typing, programming, etc. is knowledge impoverished. To capture what Anderson calls "a true functional level of the human mind," we must consider tasks that relate a person's behavior to some non-formal world.

In general, everyday and complex problem solving outside of formal domains like mathematics involves modeling the world in order to take action. Making selective observations, we construct and test alternative situation-specific models (e.g., alternative descriptions of disease processes in a particular patient), and relate them to action plans (e.g., alternative therapy processes). To understand "how our cognitive mechanisms adapt to functionally important problems," as Anderson says, we must look at problems in which the problem solver is *situated.* That is, we must study problems in which the problem solver is faced with constructing a model of the outside world, within some social setting, and relating it to the needs of some task.

More specifically, engineering problems studied in expert systems research involve modeling some system in the world (a device, a manufacturing plant, a human body, a circuit, etc.) which the person is trying to design, repair, assemble, identify, diagnosis, control, etc. (called "generic tasks" (Chandrasekaran, 1984, Clancey, 1985)). Engineering problem-solving of this type involves a modeling step to describe the world and a planning step to choose a course of action. Anderson's Lisp and geometry problems involve no world to interpret; they have no functional significance in themselves. Rather they are formal *modeling tools* that would be used in the context of some larger system-manipulating task, involving goals for doing something with this system in the world. This antecedent, mostly qualitative problem-solving, which expert systems research focuses upon, provides the analysis that leads to a theorem to prove, program to write, or equation to solve.

Qualitative models can be described on different levels (Clancey, 1986):

- the task,
- the system being modeled,
- the computational method (classification vs. construction),
- the relational network representation (e.g., prototype hierarchy, state-transition graph, procedural hierarchy), and
- the implementation in a program (rules, frames, objects, etc.).

Recurrence or "principles" include

- a vocabulary of relations for abstracting *processes in the world* (e.g., cause, progression over time, severity, location, flow-volume characteristics) and

- *cognitive processes* (often called "inference procedures") for describing complex processes in the world by explaining, predicting, or designing their behavior.

For example, routine diagnostic problem solving in medicine and sandcasting can be modeled by a common set of knowledge structures and an inference procedure (Thompson and Clancey, 1986). Routine configuration problem solving in molecular biology and construction site layout are modeled by another set of knowledge structures and inference procedure (Hayes-Roth, 1986).

Cognitive principles of this type are not necessarily explicitly stored in the brain or even articulatable by the problem solver. Rather, they are abstractions of a grammatical form that express commonalities in the behavior of individual problem solvers. These abstractions include both *kinds of patterns* experts can articulate (familiar problem-solving situations and familiar courses of action), as well as recurrent *changes in attention* and *rationales for making observations* when forming a model. An example of such recurrence is the process of "triggering" a partial model on the basis of a few observations. Triggering reflects both the cognitive ability to usefully relate a new situation to past experience and the properties of a world in which processes tend to recur. Thus, the study of recurrence of processes in cognition and the world are complementary, involving the interaction of task resources and demands.

Analysis at this level contradicts Anderson's remark that "one cannot use protocol data to analyze a skill that is already compiled." In complex problem solving such as medical diagnosis, we abstract sequences of data requests (observations made by the physician) by relating them to changes in the situation-specific model (Clancey, 1984). Moreover, if the problem solver has a model of how he reasons, as some good teachers do, we can ask for his description of the functional modeling goals that lie behind his questions (e.g., to detect errorful data, to establish temporal boundaries on the underlying cause).

While the heuristic classification (HC) model of problem solving was developed to describe expert system programs, it is also a hypothesis describing human problem solving. The HC model claims that expertise (knowledge based on experience) consists of the ability to recognize situations by abstracting specific observations and relating these system models to abstract courses of action, which are subsequently refined to the needs of the specific situation. Theories of problem solving based on such a model of experiential knowledge describe:

- a computational method (HC),
- the modeling requirements of a task (e.g., testing hypotheses, discriminating among alternative system models),
- and relevant properties of the world (e.g., nature of the recurrence in the domain, urgency, efficiency, cost of observations, importance of ordering observations).

Strikingly, problem-solving research in geometry, Lisp and Pascal programming, subtraction, algebra, etc. ignores the inherent difficulties of *modelling non-formal systems*, in which data is uncertain and incomplete, system functionality is not axiomatic, and no written calculus exists. Consequently, this research presents an impoverished view of experiential knowledge structures and inference.

In conclusion, Anderson's call to the algorithm level is reasonable, but application tasks for functionally important problems must be situated, if we are to capture cognitive principles at this level. By situated, we mean, first, that the task involves explaining and predicting events in the world in order to plan courses of action (which will in turn satisfy higher goals), and second, that the problem-solving activity is itself constrained by a social context.

In expert systems research focusing on "generic tasks," cognitive principles at the algorithm level include *representation requirements* and *inferential competence.* Representation of processes in the world involves articulating from experience models that usefully reflect the complexity of the world. Inference involves satisfying a variety of cognitive, social, and theoretical constraints when gathering information and manipulating representations, in the process of formulating adequate situation-specific models and action plans. Today's expert system shells seek to formalize representation and inference processes as principles that are useful in similar tasks in different domains.

In considering the potential of AI research for improving engineering education, we should in particular focus on the following distinctions:

- experiential knowledge vs. scientific laws
- classification vs. simulation of processes
- situated inference vs. modeling tools
- learning processes in abstracting and refining the recurrent situations and action plans of engineering practice vs. scientific methods for developing mechanistic theories.

The potential of qualitative modeling is to more adequately relate scientific theory to engineering practice, by allowing us for the first time to construct models of what engineers know from experience and how they use this knowledge in adapting designs, recognizing faults, and tuning a measurement device or manufacturing process. In particular, we can construct computer programs that are learning and problem-solving aids, built around qualitative models of physical and reasoning processes. Such models could be manipulated and inspected by students, in the course of their usual system design, diagnosis, and control activities. Most importantly, problem selection can be designed to elucidate the recurrent situations and actions familiar to experienced engineers, while revealing the exceptions and boundaries that require different approaches. Thus, we will be able to provide students with a practical context for integrating the formal theory and analysis methods, which are currently the focus of their classroom instruction, with the heretofore poorly formalized content and methods of apprenticeship learning.

References

Chandrasekaran, B. Expert systems: Matching techniques to tasks. In W. Reitman (editor), *AI Applications for Business*, pages 116-132. Ablex Publishing Corp., 1984.

Clancey, W. J. *Acquiring, representing, and evaluating a competence model of diagnosis.* HPP Memo 84-2, Stanford University, February 1984. (To appear in M. Chi, R. Glaser, and M. Farr (Eds.), *The Nature of Expertise*, in preparation.).

Clancey, W. J. Heuristic Classification. *Artificial Intelligence*, December 1985, *27*, 289-350.

Clancey, W.J. *Viewing knowledge bases as qualitative models.* KSL Report 86-27, Stanford University, 1986.

Hayes-Roth, B. *A layered environment for reasoning about actions.* HPP Memo 86-38, Stanford University, November 1986.

Thompson, T. and Clancey, W. J. A qualitative modeling shell for process diagnosis. *IEEE Software*, March 1986, *3(2)*, 6-15.